

REMARKS

Claims 1-52 are pending in the application, of which Claims 1, 11, 13, 32, 33, 41, 46 and 50-52 are independent claims. All claims stand rejected under 35 U.S.C. §§ 102(e) and/or 103(a). Claims 1, 11, 13, 32, 33, 41, 46 and 50-52 are amended and new Claim 53 is added by the present amendment. These amendments are not in acquiescence to the rejections. For the reasons stated below, it is believed that the all claims are in condition for allowance.

Claim Amendments

Claims 1, 11, 13, 32, 33, 41, 46 and 50-52 are amended to specify that the current cluster definition is stored at a single location in the shared repository, and each member node can access the current definition at this same single location. It should be understood that the single location is a logical location, which is accessible by the member nodes. The cluster definition stored at the single location can be mirrored or stored in a RAID configuration, for example. The claims are also amended to specify that the coordinator node updates the current definition at this same single location. This concept of storing, accessing and updating the cluster definition at the same single location is discussed throughout the specification and claims as originally filed. For example, as described in the application in reference to FIG. 4, “a single shared copy of the cluster definition 48 is provided in the shareable storage device 22,” and each member node can access this shared copy of the cluster definition, but only the coordinator node can update it.¹ Thus, no new matter is being introduced. Acceptance is respectfully requested.

Claims 13 and 32 are amended to specify that a member node requests a change to the cluster definition by sending a proposed change the shared repository, as similarly recited in independent Claims 1, 11, 51 and 52. Thus no new matter is being introduced.

Claim Rejections

Claims 1, 7-9, 11-13 and 14-52 were rejected under 35 U.S.C. § 102(e) based on U.S. Patent No. 6,092,213 to Lennie. Claims 2-6 and 10 were rejected under 35 U.S.C. § 103(a)

¹ See Specification at pg. 12 ll. 25 - pg. 17, ll. 2; See also Claims 1 and 5 as originally filed.

based Lennie in view of U.S. Patent No. 6,014,669 to Slaughter, U.S. Patent No. 6,003,075 to Arendt and/or U.S. Patent No. 6,243,702 to Bamford. These rejections are traversed.

In Lennie, each node maintains its own local copy of the cluster definition in the node's respective storage repository. Lennie teaches that each node has its own "safe repository" to store the cluster definition (e.g. cluster configuration data), and that a primary process executing on a coordinator node provides each of the member nodes with the cluster definition by distributing the definition to each member node. In this way, Lennie relies entirely on network connectivity to propagate the cluster definition to each member node. Like the cluster prior art,² in order to determine the cluster definition, a node is required to have network connectivity with the cluster, and the cluster definition is provided by network communication with the coordinator node.³ Consequently, a node in Lennie needs to have network connectivity before the node can be provided with the cluster definition.

By contrast, the present invention enables members of a network cluster to access a current cluster definition through a shared repository without requiring network connectivity with the cluster.⁴ A member node operating in the network cluster of the present invention does not need to have network connectivity to determine the cluster definition because each member node can access the current, updated cluster definition at a single location in the shared repository. By storing the current cluster definition at this single location, nodes no longer need network connectivity to determine the current cluster definition. Instead, member nodes only need to access the shared repository to determine the cluster definition. As a result, each member node can determine the current cluster definition without the coordinator node distributing the cluster definition to each member node. Thus, with the present invention, a member node can lose network connectivity and still determine the current cluster definition by accessing the definition at the single shared location on the shared repository.

² See e.g. Specification at pg. 12, ll. 20-24.

³ See Lennie, col. 2, ll. 44- col. 3, ll. 7.

⁴ See Specification at pg. 4, ll. 13-21.

In one embodiment, for example, cluster managers of each member node communicate with each other through disk based messaging, e.g., through a message location on the shared repository. Even though network connectivity may have been lost in the cluster, the nodes can still access the cluster definition at the shared repository and communicate with one another using their message location at the shared repository. If a node loses access to the shared repository, preferably, it removes itself from the cluster. Thus, unlike Lennie, the invention creates a network cluster that does not rely on network connectivity, but rather relies on a shared repository, which stores a single shared copy of the cluster definition.

Furthermore, when a node loses network connectivity in Lennie, there is a danger that cluster partitions, comprised of a subset of the member nodes of the cluster, may form. In the present invention, however, a node that loses connectivity can still access the cluster definition and maintain the integrity of the cluster using the shared repository.

It is further noted that in Lennie, when a node is unavailable to the cluster for a period of time, changes to the definition are stored in a master audit log. This requires a log file to be maintained, enumerating all changes to a cluster definition made while the node was unavailable. As described in the application, nodes in network clusters may occasionally be removed from the cluster for maintenance.⁵ Consequently, in Lennie, a log file could grow to a substantial size during this period of maintenance. In the present invention, however, it can be unnecessary to record information in a log file concerning nodes that have not received the current cluster definition from the coordinator node, because in the present invention, member nodes can lose network connectivity with the cluster and still access the current, updated cluster definition in the shared repository.

Thus, Lennie does not address the problems associated with a node that needs network connectivity to receive the current cluster definition, nor suggest the solutions presented in independent Claims 1, 11, 13, 32, 33, 41, 46 and 50-52. Since Lennie does not discuss the limitations of the present invention, namely, storing a current cluster definition at single location in a shared repository, accessing the definition at the single location by the member nodes, and updating the cluster definition at the single location by the coordinator node, it is believed that

⁵ See e.g. Specification at pg. 4, ll. 5-12.

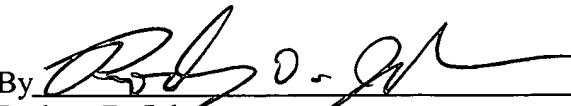
the claims are in condition for allowance. Accordingly, it is respectfully requested that the rejections to under §§ 102(e) and 103(a) be withdrawn.

CONCLUSION

In view of the above amendments and remarks, it is believed that all claims are in condition for allowance, and it is respectfully requested that the application be passed to issue. If the Examiner feels that a telephone conference would expedite prosecution of this case, the Examiner is invited to call the undersigned attorney.

Respectfully submitted,

HAMILTON, BROOK, SMITH & REYNOLDS, P.C.

By 
Rodney D. Johnson
Registration No.36,558
Telephone: (978) 341-0036
Facsimile: (978) 341-0136

Concord, MA 01742-9133

Dated: *November 17, 2003*